

AF
ZFW



**IN THE UNITED STATES PATENT AND TRADEMARK OFFICE
BEFORE THE BOARD OF PATENT APPEALS AND INTERFERENCES**

In re application of:

Carol L. Colrain, et al.

Serial No.: 10/086,782

Filed on: February 28, 2002

)
)
)
)
)
)

Confirmation No.: 8991

Examiner: Harold Dodds

Group Art Unit No.: 2168

For: **SYSTEM AND METHOD FOR PROVIDING COOPERATIVE RESOURCE
GROUPS FOR HIGH AVAILABILITY APPLICATIONS**

Mail Stop Appeal Brief – Patents
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

SUPPLEMENTAL APPEAL BRIEF

Sir:

This Supplemental Appeal Brief is submitted in support of the Notice of Appeal filed on July 7, 2006 and in response to the Notification of Non-Compliant Appeal Brief mailed on August 29, 2006.

TABLE OF CONTENTS

I.	REAL PARTY IN INTEREST
II.	RELATED APPEALS AND INTERFERENCES
III.	STATUS OF CLAIMS
IV.	STATUS OF AMENDMENTS
V.	SUMMARY OF CLAIMED SUBJECT MATTER
VI.	GROUND OF REJECTION TO BE REVIEWED ON APPEAL
VII.	ARGUMENT
A.	Teachings of <i>Gamache</i> and <i>Armentrout</i>
B.	Claims 45 and 53 Are Patentable Over <i>Gamache</i> In View of <i>Armentrout</i>
1.	<i>Gamache and Armentrout fail to teach or disclose all limitations of Claims 45 and 53</i>
2.	<i>The Examiner improperly dissected the claimed invention into discreet elements</i>
3.	<i>No proper motivation to combine Gamache and Armentrout was provided</i>
C.	Claims 46-52 and 54-60 Are Patentable Over <i>Gamache</i> In View of <i>Armentrout</i> and Further In View of <i>Dugan</i>
D.	Conclusion and Prayer For Relief
VIII.	CLAIMS APPENDIX
IX.	EVIDENCE APPENDIX
X.	RELATED PROCEEDINGS APPENDIX

I. REAL PARTY IN INTEREST

Oracle International Corporation is the real party in interest.

II. RELATED APPEALS AND INTERFERENCES

Appellants are unaware of any related appeals or interferences.

III. STATUS OF CLAIMS

Claims 1-44 have been canceled. Claims 45-60 are pending in this application, were finally rejected, and are the subject of this appeal.

IV. STATUS OF AMENDMENTS

No amendments were filed after the final Office Action.

V. SUMMARY OF CLAIMED SUBJECT MATTER

The present application contains independent Claims 45 and 53. These independent claims recite similar limitations, except in the context of a method and a computer-readable medium, respectively. Claims 45 and 53 are directed generally to an approach for providing cooperative resource groups for high availability computer applications.

Claims 45 and 53 are directed towards an approach for providing service to a service requestor in a manner that avoids the service requestor waiting for a time-out when a node providing the service fails. A service requestor uses an Internet Protocol (IP) address to address requests to a service provided by a first node within a cluster. The first node is configured to provide the service to requests addressed to the IP address (Specification, page 10, lines 5, 10, 11. In response to the first node becoming unavailable, a second node is automatically configured to respond to requests addressed to the IP address. (Specification, page 12, line 31 – page 13, line 2). After the first node becomes unavailable, the service requestor uses the IP address to address a message to the cluster related to the service (Specification, page 13, lines 2-3). In response to the message, the second node sends a response indicating an error condition (Specification, page 13, lines 2-3). (See generally Specification, page 12, line 1 – page 13, line 4).

The method and computer-readable medium of Claims 45 and 53, respectively, may exhibit numerous advantages. For example, following the approach of Claims 45 and 53, a service requestor accessing a service on a first node using an Internet Protocol (IP) address may immediately receive an indication (such as, for example, a TCP/IP error message) that the service is unavailable when the first node becomes unavailable. Advantageously, the service requestor may perform another action (such as attempting to access the service on

another node using a different IP address) faster than if the service requestor had to wait for a time-out to determine that the first node was unavailable. Other advantages and embodiments are described in further detail in the Applicants' patent application. However, the claimed invention is not required to exhibit any of these advantages. The enumerated advantages are simply examples of advantages that may be, but not need be, exhibited upon implementation of the claimed invention.

Claims 46 and 54 are, in addition to Claims 45 and 53, directed to an approach where the service requestor, upon receiving the response that indicates an error condition, identifies a second IP address in order to access the service provided by the first node (Specification, page 6, lines 3-8; page 11, lines 16-18; page 12, lines 22-25). The service requestor uses the second IP address to address a second message to the cluster related to the service provided by the first node (Specification, page 12, lines 25-27).

Claims 52 and 60 are, in addition to Claims 45 and 53, directed to an approach where the first node comprises a monitor process that is configured to detect if the first node becomes unavailable (Specification, page 10, line 14; see also "Monitors" in FIGs. 2 and 3).

VI. GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL

1. Claims 45, 46, 48, 52-54, 56, and 60 were rejected under 35 U.S.C. § 103(a) as being unpatentable over U.S. Patent No. 6,401,120 to Gamache et al. (*Gamache*) and U.S. Patent No. 6,463,457 to Armentrout et al. (*Armentrout*).

2. Claims 47, 49-51, 55, and 57-59 have been rejected under 35 U.S.C. § 103(a) as allegedly being unpatentable over *Gamache* in view of *Armentrout* and further in view of U.S. Patent Number 6,804,711 to Dugan et al. ("*Dugan*").

VII. ARGUMENTS

A. Teachings of *Gamache* and *Armentrout*

Both *Gamache* and *Armentrout* discuss vastly different techniques than that featured in the pending claims. In sharp contrast to Claims 45 and 53, *Gamache* is directed towards forming and operating a cluster from a group of available nodes. *Gamache* discusses several requirements that must be satisfied for a cluster to be formed, e.g., one node of the cluster must possess a majority of the replica members to ensure only one unique incarnation of a cluster can exist at a time (Abstract; Col. 2, lines 11-34). Importantly, while *Gamache* briefly mentions (in Col. 20, lines 10-12) that a cluster of *Gamache* may become unavailable, *Gamache* does not teach an approach for communicating, to a service requestor, that a particular node of a cluster is unavailable when that node becomes unavailable. In view of the fundamental differences between the pending claims and *Gamache*, *Gamache* does not teach numerous claim elements. Indeed, the Examiner acknowledges that *Gamache* “does not teach the use of responses requests and messages and the recognition of error conditions.” Further, the Examiner has acknowledged in an Examiner interview that *Gamache* does not disclose, teach, or suggest numerous claimed features recited in Claims 45 and 53, such as “in response to said first node becoming unavailable, automatically configuring a second node of the cluster to respond to requests addressed to said IP address” and “in response to said message, said second node of the cluster sending a response that indicates an error condition.”

Armentrout also teaches a markedly different approach than the pending claims. *Armentrout* is directed towards a platform that uses the idle computational processing power of a plurality of provider computers. A server collects tasks from client computers, schedules

and distributes the tasks to provider computers, and collects and returns results to client computers. Similarly, *Armentrout* does not teach an approach for communicating, to a service requestor, that a particular node of a cluster is unavailable when that node becomes unavailable. Further, the Examiner has acknowledged in an Examiner interview that *Armentrout* does not disclose, teach, or suggest numerous claimed features recited in Claims 45 and 53, such as “in response to said first node becoming unavailable, automatically configuring a second node of the cluster to respond to requests addressed to said IP address” and “in response to said message, said second node of the cluster sending a response that indicates an error condition.”

B. Claims 45 and 53 Are Patentable Over *Gamache* In View of *Armentrout*

To establish a prima facie case of obviousness, three basic criteria must be met. First, there must be some suggestion or motivation, either in the references themselves or in the knowledge generally available to one of ordinary skill in the art, to modify the reference or to combine reference teachings. *In re Lahu*, 747 F.2d 703, 705, 223 USPQ 1257, 1258 (Fed. Cir. 1984). Second, there must be a reasonable expectation of success. *In re Merck & Co., Inc.*, 800 F.2d 1091, 231 USPQ 375 (Fed. Cir. 1986). Finally, *the prior art reference (or references when combined) must teach or disclose all the claim limitations. In re Royka*, 490 F.2d 981, 180 USPQ 580 (CCPA 1974).

With respect to the present application, it is respectfully submitted that *Gamache* and *Armentrout*, considered alone or in combination, fail to teach or suggest all the limitations of Claims 45 and 53. Therefore, the Examiner has failed to establish a prima facie case of obviousness.

Furthermore, the Patent Office has long recognized that a proper rejection based on 35 U.S.C. § 103(a) must allege that the subject matter sought to be patented as a whole must be disclosed or taught by the prior art. MPEP § 2106, II, C states:

[W]hen evaluating the scope of a claim, every limitation in the claim must be considered. Office personnel may not dissect a claimed invention into discrete elements and then evaluate the elements in isolation. Instead, the claims as a whole must be considered. (emphasis in original)

Thus, to establish a prima facie case of obviousness as a matter of law, the claimed limitations, as a whole, must be alleged to be taught or suggested by the prior art. Merely dissecting a claimed invention into discrete elements, and then evaluating the dissected elements in isolation fails to present a prima facie case of obviousness as a matter of law.

Lastly, there is no suggestion or motivation, either in the references themselves or in the knowledge generally available to one of ordinary skill in the art, to modify the reference or to combine reference teachings.

1. Gamache and Armentrout fail to teach or disclose all limitations of Claims 45 and 53

Claims 45 and 53 recite:

a service requestor using an Internet Protocol (IP) address to address requests to a service provided by a first node within a cluster, wherein said first node is configured to provide said service to requests address to said IP address;

in response to said first node becoming unavailable, automatically configuring a second node of the cluster to respond to requests address to said IP address;

after said first node becomes unavailable, the service requestor using said IP address to address a message to said cluster related to said service; and

in response to said message, said second node of the cluster sending a response that indicates an error condition (emphasis added).

Here, the cited references, alone or in combination, fail to disclose or suggest the above-bolded features. The Examiner even acknowledged in the latest Office Action that “Gamache does not teach the use of responses[,] requests and messages and the recognition of error conditions” (Office Action, page 3). Rather than alleging that *Gamache* and *Armentrout* show meaningful portions of Claims 45 and 53, the Office Action dissects the claims into discrete portions, and then alleges that those discrete, dissected portions are shown by a cited reference.

With respect to one dissected portion of the first bolded feature, the Office Action states that “in response to said first node becoming unavailable” is shown by *Gamache* at Col. 9, lines 56-58 and Col. 20, lines 10-12, which states, *in toto*:

To create a new cluster, a system administrator runs a cluster installation utility on a system (node) that then becomes a first member of the cluster 59....For example, three replica members may be available for ordinary reliability, in which two disks will have to fail to render the **cluster unavailable**. (emphasis added).

This cited portion of *Gamache* **teaches away** from Claims 45 and 53. While *Gamache* teaches that a cluster is unavailable, Claims 45 and 53 state that the first node becomes unavailable while a second node of the cluster is able to respond to future requests to the first node. Thus, the cluster of Claims 45 and 53 is **not** unavailable. Furthermore, it is a **node** in the cluster of Claims 45 and 53 that is unavailable, not the cluster, as *Gamache* teaches. If the Examiner is correlating a disk in *Gamache* with a node in Claims 45 and 53, such a correlation is not valid. *Gamache* teaches that nodes of a cluster are separate from hard disks (e.g., see Col. 8, lines 4-8 and FIG. 3A). This is a fundamental difference between the approach of the pending claims and that of *Gamache*.

With respect to another dissected portion of the first bolded feature, the Office Action states that “automatically configuring a second node of the cluster” is shown by *Gamache* at Col. 9, lines 25-27, Col. 9, lines 62-64, and Col. 9, lines 2-4, which states, *in toto*:

To automatically failback, groups require a defined preferred owner...The **administrator then configures any resources** that are to be managed by the cluster software, possibly including other storage devices...If the failover manager 80 takes the resource offline, the group is restarted on another node in the cluster 59, known as pushing the group to another node. (emphasis added)

Again, this cited portion of *Gamache* **also teaches away** from Claims 45 and 53 for at least two reasons. First, according to the cited claim limitation, the **second node is automatically configured**, whereas *Gamache* states that **an administrator is performing a configuring operation**. Second, the claim limitation states that a **node** is configured, whereas *Gamache* states that a **resource** is configured and, in *Gamache*, a resource is different than a node in (see Col. 8, lines 4-5; Col. 5, lines 19-21, 30).

In view of the differences between what *Gamache* actually shows, and what *Gamache* is cited to show, *Gamache* cannot possibly teach or suggest the element of “in response to said first node becoming unavailable, automatically configuring a second node of the cluster.” Further, *Armentrout* is not cited to show, nor does show, this element. Thus, the combination of *Gamache* and *Armentrout* also fails to teach or suggest this element as well.

With respect to a dissected portion of the second bolded feature, the Office Action states that “after said first node becomes unavailable” is shown by *Gamache* at Col. 9, lines 56-58, Col. 20, lines 10-12, which states, *in toto*:

To create a new cluster, a system administrator runs a cluster installation utility on a system (node) that then becomes a first member of the cluster 59...For example, three replica members may be available for ordinary reliability, in which **two disks will have to fail** to render the cluster unavailable. (emphasis added)

According to the cited claim limitation, the first node becomes unavailable. However, the Examiner cited *Gamache* for teaching (1) that a node becomes a member of a cluster and (2) that two disks must fail to render a cluster unavailable. However, as stated previously, the disks and nodes of *Gamache* are **not equivalent** (see Col. 8, lines 4-8 and FIG. 3A). Therefore, it would not be obvious for someone skilled in the art to read the above cited portion of *Gamache* and determine that the first node becomes unavailable.

Assuming, *arguendo*, that *Gamache* disclosed “first node becomes unavailable”, *Gamache* would still fail to teach or suggest “the service requestor using said IP address to address a message to said cluster related to said cluster” **after** the “first node becomes unavailable”, as recited by Claims 45 and 53. Because *Gamache* fails to disclose that a service requestor uses an IP address to address a message to a cluster **after** the first node becomes unavailable, the combination of *Gamache* and *Armentrout* also fails to disclose this feature.

Furthermore, *Gamache* fails to teach or suggest that a service requestor addresses a message to a cluster, much less “the service requestor using said IP address to address a message to said cluster”, as recited in Claims 45 and 53. Although the Examiner admitted that “*Gamache* does not teach the use of responses[,] requests and messages and the recognition of error conditions”, the Office Action states that this feature is shown by *Gamache* at col. 7, lines 30-32, col. 5, lines 39-43, col. 7, lines 19-25, col. 5, line 66 – col. 6, line 2, which states, *in toto*:

With GLUP, a node (e.g., 60₂) wishing to send an update to other nodes first sends a request to the locker node 60₁... For example, a group may include an application that depends on a network name, which in turn depends on an Internet Protocol (IP) address, all of which are collected in a single group... The global update protocol (GLUP) is used by the global update manager 90 to broadcast updates to each node

60.sub.1 -60.sub.n in the cluster 59. GLUP generally comprises a standard global update message format, state information maintained in each node, and a set of rules that specify how global update should be processed and what steps should be taken when failures occur...Through the communications manager 72, the cluster server 66 is essentially in constant communications with the other nodes 60₁ – 60_n of the cluster 59.

The above cited portions of *Gamache* disclose that only a node, update manager, or cluster server sends messages to nodes in the cluster. No where does *Gamache* teach that a **service requestor** addresses messages to a node in the cluster, as recited in Claims 45 and 53.

Gamache fails to teach or suggest the above dissected portion of the claim. Therefore, *Gamache* and *Armentrout*, individually or in combination, cannot possibly teach or suggest “after said first node becomes unavailable, automatically configuring a second node of the cluster to respond to requests addressed to said IP address.”

Based on the foregoing, *Gamache* and *Armentrout*, in combination, fail to teach numerous features of Claims 45 and 53. Fundamentally, their combination fails to illustrate the concept of a service requestor addressing messages to an IP address where a node (that is configured to provide a service to requestors that address messages to the node) becomes unavailable and another node responds in place of the first node, much less that the other node is automatically configured to respond to requests addressed to the IP address. It is therefore respectfully submitted that Claims 45 and 53 are patentable over the cited references.

2. The Examiner improperly dissected the claimed invention into discreet elements

35 U.S.C. § 103(a) states:

“A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the

subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains” (emphasis added).

Thus, as a matter of law, to present a prima facie case of obviousness under 35 U.S.C. § 103(a), the differences between the subject matter sought to be patented and the prior art must be such that the subject matter sought to be patented as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains.

The Patent Office has long recognized that a proper rejection based on 35 U.S.C. § 103(a) must allege that the subject matter sought to be patented as a whole must be disclosed or taught by the prior art. MPEP § 2106, II, C states:

[W]hen evaluating the scope of a claim, every limitation in the claim must be considered. Office personnel may not dissect a claimed invention into discrete elements and then evaluate the elements in isolation. Instead, the claims as a whole must be considered. (emphasis in original)

Thus, to establish a prima facie case of obviousness as a matter of law, the claimed limitations, as a whole, must be alleged to be taught or suggested by the prior art. **Merely dissecting a claimed invention into discrete elements, and then evaluating the dissected elements in isolation fails to present a prime facie case of obviousness as a matter of law.**

In violation of MPEP § 2106, II, C and 35 U.S.C. § 103(a), the Examiner not only divided each claim into its constituent elements, but also finely dissected each constituent element into a set of short phrases and sentence fragments. The Examiner then pointed out how each individual fragment corresponds to a similar fragment in any one of a handful of references. **The fragment-to-prior-art correlation appears to have been made without**

any consideration as to the relationship between the fragments, the meaning of the elements as a whole, and the meaning of the claim as a whole. If rejecting a claim in this manner was proper, nearly every claim may be rejected since most claims use words or phrases found in other documents.

None of the cited references have been cited to disclose, teach, or suggest meaningful portions of subject matter recited in the independent claims. Instead, the Examiner dissected the claimed invention into discrete elements and then evaluated the dissected elements **in isolation**, rather than considering the elements as a whole.

For example, no cited reference is alleged to show the subject matter of “in response to said first node becoming unavailable, automatically configuring a second node of the cluster to respond to requests associated with said IP address.” Instead:

- One portion of *Gamanche* is cited to show “in response to said first node becoming unavailable”,
- another portion of *Gamanche* is cited to show “automatically configuring a second node of the cluster”,
- *Armentrout* is cited to show “to respond to requests”, and
- another portion of *Gamanche* is cited to show “associated with said IP address.”

The above dissected portions have been evaluated in isolation, and are divorced from their actual meaning given their placement in the claim. For example, this element recites the performance of an action (automatically configuring a second node) in response to a condition being satisfied (the first node becoming unavailable). No reference shows, or is cited to show, the performance of this action in response to this condition being satisfied.

As another example of dissecting the claimed invention into discrete elements and then evaluating the dissected elements in isolation, no cited reference is alleged to show the

subject matter of “in response to said message, said second node of the cluster sending a response that indicates an error condition.” Instead:

- one portion of *Armentrout* is cited to show “in response to said message”,
- one portion of *Gamache* is cited to show “said second node of the cluster”,
- another portion of *Armentrout* is cited to show “sending a response that indicates an error condition.”

One result of citing portions of one or more references where the cited portions have no relation to each other is that a correlated feature in a reference is not consistent throughout the rejection. In citing both *Gamache* and *Armentrout*, the Examiner failed to keep consistent the correlated features within *Gamache* and between *Gamache* and *Armentrout*.

For example, in the last Office Action reply mailed February 14, 2006, the Applicants stated:

Claim 45 also features the element of “in response to said message, said second node of the cluster sending a response that indicates an error condition.” Neither *Gamache* nor *Armentrout* suggest the performance of this element. The Office Action acknowledges that *Gamache* “does not teach the user of responses requests and messages and the recognition of error conditions.” Despite this acknowledgement of the shortcoming of *Gamache*’s teachings, the Office Action relies upon *Gamache* to show the fragment of “said second node of the cluster,” and relies upon *Armentrout* (Col. 23, lines 64-67, Col. 17, lines 62-64, and Col. 24, line 2) to show the remainder of this element.

However, the cited portions of *Armentrout* lack (a) any suggestion that the second node of the cluster, as claimed, is performing any actions, and (b) any suggestion of sending a response that indicates an error condition in response to a message sent to the cluster by the service requestor. (emphasis in original)

The Examiner responded to this argument with the following:

The Applicants err when they suggest that *Armentrout* does not teach the use of nodes with clusters. In fact, *Armentrout* associates clusters with nodes as follows:

“In operation, the present invention allows the client to specify job and task parameters via the client interface. Using the client interface, the Client selects desired attributes of nodes in a virtual cluster...” at col. 13, lines 4-7.

If this *Armentrout* citation were substituted for the *Gamache* citing for this limitation, the prior art rendering obvious this limitation would consist entirely of *Armentrout* citations.

By admitting that *Armentrout* alone could have been used to render this limitation obvious, the Examiner has provided another example in how the claimed invention was dissected into discrete elements. The Examiner cited the above portion of *Armentrout* for teaching “second node of the cluster.” Thus, the above-quoted response by the Examiner indicates that *Armentrout* teaches “in response to said message, a second node of the cluster sending a response that indicates an error condition”. However, although *Armentrout* uses the terms “in response to”, “node”, “cluster”, “message” and “error”, *Armentrout* fails to teach or suggest that a **second node of the cluster sends a response**, much less that the response is in response to a message. Furthermore, even if *Gamache* disclosed the first three steps of Claims 45 and 53, *Armentrout* gives no teaching or suggestion **that the message that initiates the response is a message that** (1) came from a service requestor, (2) was sent to a cluster, and (3) was addressed using an IP address by which a first node is configured to provide a service request.

As stated in MPEP § 2106, II, C, **Office personnel may not dissect a claimed invention into discrete elements and then evaluate the elements in isolation**. It is respectfully submitted that, rather than alleging that the subject matter of “in response to said first node becoming unavailable, automatically configuring a second node of the cluster to respond to requests associated with said IP address” is taught by any cited reference, the Examiner impermissibly dissected the claimed invention into discrete elements and then evaluated those elements in isolation. The same applies to the other two steps of Claims 45 and 53 bolded above.

Accordingly, it is submitted that the Examiner failed to present a prima facie case of obviousness under 35 U.S.C. § 103(a) as a matter of law. As a consequence, it is respectfully submitted that the rejection made to Claims 43-60 under 35 U.S.C. § 103(a) may not be properly maintained.

3. *No proper motivation to combine Gamache and Armentrout was provided*

Gamache and *Armentrout* have not been properly combined. The Office Action states:

It would have been obvious to one of ordinary skill at the time of the invention to combine *Armentrout* with *Gamache* to provide responses to requests and messages in order to use standard communication protocol between members of a network and to promote user acceptance of the system. Likewise, it would have been obvious to one of ordinary skill at the time of the invention to combine *Armentrout* with *Gamache* to detect error conditions and to permit appropriate action to be taken to correct the problem or take another corrective action and thus provide a more stable system. *Gamache* and *Armentrout* teach the user of related systems. They teach the use of computers, the use of databases, the use of networks, the use of nodes, the use of clusters, the use of services, the use of resources, the use of applications, the monitoring of status, the use of paths, and the detection of failures.

However, notwithstanding the fact that the Examiner acknowledged that neither *Gamache* nor *Armentrout* disclose numerous elements of Claim 45, the Applicants respectfully submit that there is nothing in either *Gamache* or *Armentrout* that teaches or suggests combining their respective teachings. Indeed, the Examiner has not recited any reference to support the existence of a motivation to combine the references but instead broadly concludes that such a combination would have been obvious without support or justification.

As stated in the Federal Circuit decision *In re Dembiczak*, 50 USPQ.2d 1617 (Fed. Cir. 1999), (citing *Gore v. Garlock*, 220 USPQ 303, 313 (Fed. Cir. 1983)), “it is very easy to fall victim to the insidious effect of the hindsight syndrome where that which only the inventor taught is used against its teacher.” *Id.* The Federal Circuit stated in *Dembiczak* “that the best defense against subtle but powerful attraction of a hindsight-based obviousness analysis is rigorous application of the requirement for a showing of the teaching or suggestion to combine prior art references.” *Id.* Thus, the Federal Circuit explains that a

proper obviousness analysis requires “*particular factual findings* regarding the locus of the suggestion, teaching, or motivation to combine prior art references.” *Id.* (emphasis added).

In particular, the Federal Circuit states:

“We have noted that evidence of a suggestion, teaching, or motivation to combine may flow from the prior art references themselves, the knowledge of one of ordinary skill in the art, or, in some cases, from the nature of the problem to be solved...although ‘the suggestion more often comes from the teachings of the pertinent references’...The range of sources available, however, does *not diminish the requirement for actual evidence*. That is, the *showing must be clear and particular*...Broad conclusory statements regarding the teaching of multiple references, standing alone, are not ‘evidence.’” *Id.* (emphasis added; internal citations omitted).

Neither *Gamache* nor *Armentrout* show any suggestion, teaching, or motivation to combine their teachings, nor did the Examiner provide a “clear and particular” showing of the suggestion, teaching, or motivation to combine their teachings. Both *Gamache* and *Armentrout* are directed towards drastically different subject matter, namely providing consistent cluster operation data and utilizing idle computational processing power. The Examiner’s unsupported assertion that the references both use “resources,” “nodes,” “computers,” “databases,” “networks,” “clusters,” “services,” “applications,” “monitoring of status,” “paths,” and “the detection of failures” does not provide a motivation for combination, as these features are so common, **any reference may be said to be combined with another reference in this fashion based on a key word search of nouns** which both references recite therein, irrespective of the teachings of the references. The assertion also directly contradicts the Examiner’s earlier statement that “*Gamache* does not teach...the recognition of error conditions.”

No portion of any reference is cited by the Examiner to provide a motivation for combination of the references. In fact, the only motivation provided by the Examiner is the

hindsight observation that by combining features of those references, one may achieve the benefits achieved from the invention as described and claimed in the application. It is respectfully submitted that such a hindsight observation is not consistent with the Federal Circuit's requirement for "particular factual findings."

C. Claims 46-52 and 54-60 Are Patentable Over *Gamache* In View of *Armentrout* and Further In View of *Dugan*

Claims 46-52 and 54-60 are dependant claims that depend on an independent claim that is discussed above. Because each of the dependant claims includes the limitations of claims upon which they depend, the dependant claims are patentable for at least those reasons the claims upon which the dependant claims depend are patentable. In addition, the dependent claims introduce additional limitations that independently render them patentable.

For example, Claims 46 and 54 recite:

upon receiving said response, the **service requestor identifying a second IP address to access said service**; and

the service requestor using said second IP address to address a second message to said cluster related to said service. (emphasis added)

The Examiner cited Armentrout for teaching "upon receiving said response". The Examiner then cited Col. 7, lines 30-32 and col. 5, lines 37-43 as teaching "the service requestor identifying a second IP address to access said service", which states, *in toto*:

With GLUP, a node (e.g., 60₂) wishing to send an update to other nodes first sends a request to the locker node 60₁... Usually a group contains all of the elements needed to run a specific application, and for client systems to connect to the service provided by the application. For example, a group may include an application that depends on a network name, which in turn depends on an Internet Protocol (IP) address, all of which are collected in a single group.

Claims 46 and 54 require that the service requestor identifies a second IP address to access the service **upon receiving a response** that indicates an error condition. *Gamache* fails to

teach or even suggest this, much less teach or suggest that the service requestor identifies a second IP address **after using a first IP address**.

As another example of dependent claims that introduce additional limitations that independently render them patentable, Claims 52 and 60 state “wherein said first node comprises a monitor process, and wherein said monitor process is configured to detect if said first node becoming unavailable.” The Examiner cited six different sections of *Gamache* for teaching this feature, which state, *in toto*:

To create a new cluster, a system administrator runs a cluster installation utility on a system (node) that then becomes a **first member** of the cluster 59. For a new cluster 59... A **resource monitor** 82 runs in one or more processes that may be part of the cluster service 66, but are shown herein as being separate from the cluster service 66 and communicating therewith via RPC or the like... The administrator then **configures any resources** that are to be managed by the cluster software, possibly including other storage devices... A first possible outcome to the read request is that the **read will fail** (as detected at step 904) because another node (e.g., 60₂) has previously placed (and not released) a reservation on the quorum member 58₁ ... To create a new cluster, a system administrator runs a cluster installation utility on a system (node) that then becomes a first member of the cluster 59. For a new cluster 59, a total replica set 106 of replica members is created, each member including a database (e.g., 100₁, FIG. 4) having an identical copy of the initial cluster operational data added thereto... For example, three replica members may be available for ordinary reliability, in which two disks will have to fail to render the **cluster unavailable**. (emphasis added)

Although *Gamache* uses the terms “node”, “monitor”, “configures”, “detect”, and “unavailable” *Gamache* fails to teach or suggest that **the monitor process is configured to detect the node becoming unavailable**. Again *Gamache* teaches that an administrator configures a resource, which is not a monitor process. Furthermore, *Gamache* fails to teach that the monitor process detects anything relating to actions of a node, much less that the node is unavailable. *Gamache* simply states that the resource monitor 82 **monitors the health of resources**, which, as discussed above, **are not nodes** in the cluster.

D. Conclusion and Prayer For Relief


Based on the foregoing, it is respectfully submitted that the rejection of Claims 45, 46, 48, 52-54, 56, and 60 under 35 U.S.C. § 103 as being unpatentable over *Gamache* in view of *Armentrout* lacks the requisite factual and legal bases. Appellants therefore respectfully request that the Honorable Board reverse the rejection of Claims 45, 46, 48, 52-54, 56, and 60 under 35 U.S.C. § 103 over *Gamache* in view of *Armentrout*.

It is also respectfully submitted that the rejection of Claims 47, 49-51, 55, and 57-59 under 35 U.S.C. § 103 as being unpatentable over *Gamache* in view of *Armentrout* and further in view of *Dugan* lacks the requisite factual and legal bases. Appellants therefore respectfully request that the Honorable Board reverse the rejection of Claims 47, 49-51, 55, and 57-59 under 35 U.S.C. § 103 over *Gamache* in view of *Armentrout* and further in view of *Dugan*.

Respectfully submitted,

HICKMAN PALERMO TRUONG & BECKER LLP

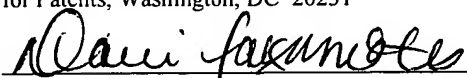
Date: December 4, 2006


Daniel D. Ledesma
Reg. No. 57,181

2055 Gateway Place, Suite 550
San Jose, California 95110-1089
Tel: (408) 414-1229
Fax: (408) 414-1076

CERTIFICATE OF MAILING

I hereby certify that this correspondence is being deposited with the United States Postal Service as first class mail in an envelope addressed to: Mail Stop Appeal Brief – Patents, Commissioner for Patents, Washington, DC 20231

on December 4, 2006 by 
Darci Sakamoto

XIII. CLAIMS APPENDIX

45. A machine-implemented method, comprising the steps of:
a service requestor using an Internet Protocol (IP) address to address requests to a service provided by a first node within a cluster, wherein said first node is configured to provide said service to requests addressed to said IP address;
in response to said first node becoming unavailable, automatically configuring a second node of the cluster to respond to requests addressed to said IP address;
after said first node becomes unavailable, the service requestor using said IP address to address a message to said cluster related to said service; and
in response to said message, said second node of the cluster sending a response that indicates an error condition.
46. The method of Claim 45, further comprising the steps of:
upon receiving said response, the service requestor identifying a second IP address to access said service; and
the service requestor using said second IP address to address a second message to said cluster related to said service.
47. The method of Claim 45, further comprising the step of:
storing, at the first node, information identifying one or more nodes of the cluster as being standby nodes, wherein each of the one or more standby nodes may be instructed to provide the service if the first node becomes unavailable.
48. The method of Claim 45, further comprising the step of:
in response to said first node becoming unavailable, determining if said first node is configured to allow the service to be provided by another node of the cluster.

49. The method of Claim 48, further comprising the step of:
in response to determining said first node is configured to allow the service to be provided by another node of the cluster, determining a standby node of the cluster to perform the service; and
instructing the standby node to perform the service.
50. The method of Claim 45, further comprising the steps of:
in response to said first node becoming unavailable, instructing a standby node of the cluster to perform the service;
determining if the plurality of services provided by the standby node may be provided by another node of the cluster; and
if the plurality of services provided by the standby node may not be provided by another node of the cluster, configuring the standby node to disallow the plurality of services to be provided by another node of the cluster.
51. The method of Claim 50, further comprising the step of:
in response to configuring the standby node to disallow the plurality of services to be provided by another node of the cluster, issuing an alert to a user.
52. The method of Claim 45, wherein said first node comprises a monitor process, and wherein said monitor process is configured to detect if said first node becoming unavailable.

53. A machine-readable medium carrying one or more sequences of instructions, wherein execution of the one or more sequences of instructions by one or more processors causes the one or more processors to perform the steps of, comprising the steps of: a service requestor using an Internet Protocol (IP) address to address requests to a service provided by a first node within a cluster, wherein said first node is configured to provide said service to requests addressed to said IP address; in response to said first node becoming unavailable, automatically configuring a second node of the cluster to respond to requests addressed to said IP address; after said first node becomes unavailable, the service requestor using said IP address to address a message to said cluster related to said service; and in response to said message, said second node of the cluster sending a response that indicates an error condition.
54. The machine-readable medium of Claim 53, wherein execution of the one or more sequences of instructions by the one or more processors causes the one or more processors to perform the steps of: upon receiving said response, the service requestor identifying a second IP address to access said service; and the service requestor using said second IP address to address a second message to said cluster related to said service.
55. The machine-readable medium of Claim 53, wherein execution of the one or more sequences of instructions by the one or more processors causes the one or more processors to perform the step of: storing, at the first node, information identifying one or more nodes of the cluster as being standby nodes, wherein each of the one or more standby nodes may be instructed to provide the service if the first node becomes unavailable.

56. The machine-readable medium of Claim 53, wherein execution of the one or more sequences of instructions by the one or more processors causes the one or more processors to perform the step of:
in response to said first node becoming unavailable, determining if said first node is configured to allow the service to be provided by another node of the cluster.
57. The machine-readable medium of Claim 56, wherein execution of the one or more sequences of instructions by the one or more processors causes the one or more processors to perform the step of:
in response to determining said first node is configured to allow the service to be provided by another node of the cluster, determining a standby node of the cluster to perform the service; and
instructing the standby node to perform the service.
58. The machine-readable medium of Claim 53, wherein execution of the one or more sequences of instructions by the one or more processors causes the one or more processors to perform the steps of:
in response to said first node becoming unavailable, instructing a standby node of the cluster to perform the service;
determining if the plurality of services provided by the standby node may be provided by another node of the cluster; and
if the plurality of services provided by the standby node may not be provided by another node of the cluster, configuring the standby node to disallow the plurality of services to be provided by another node of the cluster.
59. The machine-readable medium of Claim 58, wherein execution of the one or more sequences of instructions by the one or more processors causes the one or more processors to perform the step of:
in response to configuring the standby node to disallow the plurality of services to be provided by another node of the cluster, issuing an alert to a user.

60. The machine-readable medium of Claim 53, wherein said first node comprises a monitor process, and wherein said monitor process is configured to detect if said first node becoming unavailable.

IX. EVIDENCE APPENDIX

None

X. RELATED PROCEEDINGS INDEX

None